

# Fully Parallel Particle Learning for GPGPUs and Other Parallel Devices

Kenichiro McAlinn\*      Hiroaki Katsura<sup>†</sup>      Teruo Nakatsuma<sup>‡</sup>

First draft: December, 2012

## Abstract

We developed a novel parallel algorithm for particle filtering (and learning) which is specifically designed for GPUs (graphics processing units) or similar parallel computing devices. In our new algorithm, a full cycle of particle filtering (computing the value of the likelihood for each particle, constructing the cumulative distribution function (CDF) for resampling, resampling the particles with the CDF, and propagating new particles for the next cycle) can be executed in a massively parallel manner. One of the advantages of our algorithm is that every single numerical computation or memory access related to the particle filtering is executed solely inside the GPU, and no data transfer between the GPU's device memory and the CPU's host memory occurs unless it is under the absolute necessity of moving generated particles into the host memory for further data processing, so that it can circumvent the limited memory bandwidth between the GPU and the CPU. To demonstrate the advantage of our parallel algorithm, we conducted a Monte Carlo experiment in which we applied the parallel algorithm as well as conventional sequential algorithms for estimation of a simple state space model via particle learning, and compared them in terms of execution time. The results showed that the parallel algorithm was far superior to the sequential algorithm.

**Keywords:** Bayesian learning, GPGPU, parallel computing, particle filtering, state space model.

---

\*Graduate School of Economics, Keio University, 2-15-45 Mita, Minato-ku, Tokyo, Japan

<sup>†</sup>Graduate School of Economics, Keio University, 2-15-45 Mita, Minato-ku, Tokyo, Japan

<sup>‡</sup>Faculty of Economics, Keio University, Keio University, 2-15-45 Mita, Minato-ku, Tokyo, Japan

# 1 Introduction

The state space model (SSM) has been one of the indispensable tools for time series analysis and optimal control for decades. Although the archetypal SSM is linear and Gaussian, the literature of more general non-linear and non-Gaussian SSMs have been rapidly growing in the last two decades. For lack of an analytically tractable way to estimate the general SSM, numerous approximation methods have been proposed. Among them, arguably the most widely applied method is particle filtering (Gordon et al. (1993), Kitagawa (1996)). Particle filtering is a type of sequential Monte Carlo method in which the integrals we need to evaluate for filtering are approximated by the Monte Carlo integration. To improve numerical accuracy and stability of the particle filtering algorithm, various extensions such as the auxiliary particle filter (Pitt and Shephard (1999)) have been proposed, and still actively studied by many researchers. For SSMs with unknown parameters, Kitagawa (1998) proposed a self-organizing state space modeling approach in which the unknown parameters are regarded as a subset of the state variables and the joint posterior distribution of the parameters and the state variables is evaluated with a particle filtering algorithm. Other particle filtering methods that can simultaneously estimate parameters have been proposed by Liu and West (2001), Storvik (2002), Fearnhead (2002), Polson et al. (2008), Johannes and Polson (2008), Johannes et al. (2008), Carvalho et al. (2010), just to name a few. These particle filtering methods that estimate state variables and parameters simultaneously are often called particle learning methods in the literature. Although the effectiveness of particle filtering methods have been proven through many different applications (see Zou and Chakrabarty (2007), Mihaylova et al. (2007), Chai and Yang (2007), Montemerlo et al. (2003), Dukic et al. (2009), and Lopes and Tsay (2011) among others), it is offset by the fact that it is a time-consuming technique. Some practitioners still shy away from using it in their applications because of this.

This attitude toward particle filtering would be changed by the latest technology: parallel computing. As we will discuss in Section 2, some parts of the particle filtering procedure are ready to be executed simultaneously on many processors in a parallel computing environment. In light of inexpensive parallel processing devices such as GPGPUs<sup>1</sup> (general purpose

---

<sup>1</sup>A high-performance GPU (graphics processing unit) was originally developed for displaying high-resolution 2D/3D graphics necessary in video games and computer-aided design. Because a GPU is designed with a massive

graphics processing units) available to the general public, more and more researchers start to jump on the bandwagon of parallel computing. Lee et al. (2010) reviewed general attempts at parallelization of Bayesian estimation methods. Durham and Geweke (2011) developed a sequential Monte Carlo method designed for GPU computing and applied it to complex non-linear dynamic models, which are numerically intractable even for the Markov chain Monte Carlo method. As for parallelization of particle filtering, a few researches (see Montemayor et al. (2004), Maskell et al. (2006), Hendeby et al. (2007) for example) have been reported, though the field is still in a very early stage. To the best of the authors' knowledge, Hendeby et al. (2010) is the only successful implementation of the particle filtering algorithm on the GPU. Their implementation, however, depends on device-specific functionalities and its resampling algorithm is not an exact one.

In developing parallel algorithms designed for the GPU, there are a few bottlenecks one should avoid. First, processing sequential algorithms on the GPU can be inefficient because of the GPU's device memory architecture. Roughly speaking, a GPU has two types of memory: memory assigned to each core and memory shared by all cores. Access to the core-linked memory is fast while access to the shared memory takes more time. Ideally, one should try as much to keep all calculations on each core without any communications among cores. The second bottleneck is that it is time-consuming to transfer memory between the host memory, which the CPU uses, and the device memory, which the GPU uses. In other words the bandwidth between the GPU's device memory and the CPU's host memory is very narrow. Thus an ideal parallel algorithm for the GPU would be to calculate everything within the GPU, preferably within each core (without inter-core communications).

With these bottlenecks in mind, we have developed a new parallel algorithm that computes the full cycle of the particle filtering algorithm in a massively and fully parallel manner, from computing the likelihood for each particle, constructing the cumulative distribution function (CDF) for resampling, resampling the particles with the CDF, and propagating new particles

---

number of processor cores to conduct single-instruction multiple-data (SIMD) processing, it has been regarded as an attractive platform of parallel computing and researchers started to use it for high-performance computing. As GPU manufacturers try to take advantage of this opportunity, it has evolved into a more computation-oriented device called GPGPU. Nowadays almost all GPUs have more or less capabilities for parallel computing, so the distinction between GPUs and GPGPUs are blurred.

for the next cycle. By keeping all of our computations within the GPU and avoiding all memory transfer between the GPU and the CPU during the execution of the particle filtering algorithm. In this way, we exploit the great benefits of parallel computing on the GPU while avoiding its short comings. Additionally, since our parallel algorithm does not utilize any device-specific functionalities, it can be easily implemented on other parallel computing devices including cloud computing systems.

In order to compare our new parallel algorithm with conventional sequential algorithms, we conducted a Monte Carlo experiment in which we applied the competing particle learning algorithms to estimate a simple state space model (stochastic trend with noise model) and recorded the execution time of each algorithm. The results showed that our parallel algorithm on the GPU is far superior to the conventional sequential algorithm on the CPU by around 30x to 200x.

The organization of this paper is as follows. In Section 2, we briefly review state space models and particle filtering and learning. In Section 3, we describe how to implement a fully parallelized particle filtering algorithm, in particular how to parallelize the CDF construction step and the resampling step. In Section 4, we report the results of our Monte Carlo experiment and discuss their implications. In Section 5, we state our conclusion.

## 2 State Space Models and Particle Filtering

A general form of SSM is given by

$$\begin{cases} y_t \sim p(y_t|x_t) \\ x_t \sim p(x_t|x_{t-1}) \end{cases} \quad (1)$$

where  $p(y_t|x_t)$  stands for the conditional distribution or density of observation  $y_t$  given unobservable  $x_t$  and  $p(x_t|x_{t-1})$  stands for the conditional distribution or density of  $x_t$  given  $x_{t-1}$ , which is the previous realization of  $x_t$  itself. In the literature of SSM, unobservable  $x_t$ , which dictates the stochastic process of  $y_t$ , is called the state variable.

Time series data analysis with the SSM is centered on how to dig up hidden structures of the state variable out of the observations  $\{y_t\}_{t=1}^T$ . In particular, the key questions in applications of SSM are (i) how to estimate the current unobservable  $x_t$ , (ii) how to predict the future

state variables, and (iii) how to infer about the past state variables with the data currently available. These aspects of state space modeling are called *filtering*, *prediction*, and *smoothing*, respectively.

The filtering procedure, which is the main concern in our study, is given by the sequential Bayes filter:

$$p(x_t|y_{1:t-1}) = \int p(x_t|x_{t-1})p(x_{t-1}|y_{1:t-1})dx_{t-1}, \quad (2)$$

$$p(x_t|y_{1:t}) = \frac{p(y_t|x_t)p(x_t|y_{1:t-1})}{\int p(y_t|x_t)p(x_t|y_{1:t-1})dx_t}, \quad (3)$$

where  $y_{1:t} = \{y_1, \dots, y_t\}$  ( $t = 1, \dots, T$ ) and  $p(x_t|y_{1:t})$  is the conditional density of the state variable  $x_t$  given  $y_{1:t}$ . In essence, equation (3) is the well-known Bayes rule to update the conditional density of  $x_t$  while equation (2) is the one-period-ahead predictive density of  $x_t$  given the past observations  $y_{1:t-1}$ . By applying (2) and (3) repeatedly, one keeps the conditional density  $p(x_t|y_{1:t})$  updated as a new observation comes in.

In general, a closed-form of neither (2) nor (3) is available, except for the linear Gaussian case where we can use the Kalman filter (Kalman (1960)). See West and Harrison (1997) on detailed accounts about the linear Gaussian SSM. To deal with this difficulty, we apply particle filtering, in which we approximate the integrals in (2) and (3) with *particles*, a random sample of the state variables generated from either the conditional density  $p(x_t|y_{1:t})$  or the predictive density  $p(x_t|y_{1:t-1})$ . Let  $\{x_t^{(i)}\}_{i=1}^N$  denote  $N$  particles generated from  $p(x_t|y_{1:t-1})$ . We can approximate  $p(x_t|y_{1:t-1})$  by

$$p(x_t|y_{1:t-1}) \simeq \frac{1}{N} \sum_{i=1}^N \delta(x_t - x_t^{(i)}) \quad (4)$$

where  $\delta(\cdot)$  is the Dirac delta. Then the filtering equation (3) is approximated by

$$p(x_t|y_{1:t}) \simeq \frac{p(y_t|x_t) \frac{1}{N} \sum_{i=1}^N \delta(x_t - x_t^{(i)})}{\int p(y_t|x_t) \frac{1}{N} \sum_{i=1}^N \delta(x_t - x_t^{(i)}) dx_t} = \sum_{i=1}^N w_t^{(i)} \delta(x_t - x_t^{(i)}), \quad (5)$$

$$w_t^{(i)} = \frac{p(y_t|x_t^{(i)})}{\sum_{i=1}^N p(y_t|x_t^{(i)})}.$$

(5) implies that the conditional density  $p(x_t|y_{1:t})$  is discretized on particles  $\{x_t^{(i)}\}_{i=1}^N$  with probabilities  $\{w_t^{(i)}\}_{i=1}^N$ . Therefore we can obtain a sample of  $x_t$ ,  $\{\tilde{x}_t^{(i)}\}_{i=1}^N$ , from  $p(x_t|y_{1:t})$  by drawing each  $\tilde{x}_t^{(i)}$  out of  $\{x_t^{(i)}\}_{i=1}^N$  with probabilities  $\{w_t^{(i)}\}_{i=1}^N$  when the approximation (5) is sufficiently accurate. This procedure is called *resampling*. In reverse, if we have  $N$  particles  $\{\tilde{x}_{t-1}^{(i)}\}_{i=1}^N$  generated from  $p(x_{t-1}|y_{1:t-1})$ , we can approximate  $p(x_t|y_{1:t-1})$  by

$$p(x_t|y_{1:t-1}) \simeq \int p(x_t|x_{t-1}) \frac{1}{N} \sum_{i=1}^N \delta(x_{t-1} - \tilde{x}_{t-1}^{(i)}) dx_{t-1} = \frac{1}{N} \sum_{i=1}^N p(x_t|\tilde{x}_{t-1}^{(i)}). \quad (6)$$

Then (6) implies that we can obtain a sample of  $x_{t+1}$ ,  $\{x_{t+1}^{(i)}\}_{i=1}^N$ , from  $p(x_{t+1}|y_{1:t})$  by generating each  $x_{t+1}^{(i)}$  from  $p(x_{t+1}|\tilde{x}_t^{(i)})$ , which is called *propagation*. Hence we can mimic the sequential Bayes filter by repeating the *propagation* equation (6) and the *resampling* equation (5) for  $t = 1, 2, 3, \dots$ . This is the basic principle of particle filtering. The formal representation of the particle filtering algorithm is given as follows.

#### ALGORITHM: PARTICLE FILTERING

**Step 0:** Set the starting values of  $N$  particles  $\{\tilde{x}_0^{(i)}\}_{i=1}^N$ .

**Step 1:** Propagate  $x_t^{(i)}$  from  $p(x_t|\tilde{x}_{t-1}^{(i)})$ , ( $i = 1, \dots, N$ ).

**Step 2:** Compute weight  $w_t^{(i)} \propto p(y_t|x_t^{(i)})$ .

**Step 3:** Resample  $\{\tilde{x}_t^{(i)}\}_{i=1}^N$  from  $\{x_t^{(i)}\}_{i=1}^N$  with weight  $w_t^{(i)}$ .

When a state space model depends on unknown parameters  $\theta$ ,

$$\begin{cases} y_t \sim p(y_t|x_t, \theta) \\ x_t \sim p(x_t|x_{t-1}, \theta) \end{cases} \quad (7)$$

we need to evaluate the posterior distribution  $p(\theta|y_{1:t})$  given the observations  $y_{1:t}$ . In the framework of particle filtering,  $p(\theta|y_{1:t})$  is sequentially updated as a new observation arrives, which is called particle learning. The particle learning algorithm is defined as follows. Let  $\{z_t^{(i)} = (x_t^{(i)}, \theta_t^{(i)})\}_{i=1}^N$  and  $\{\tilde{z}_t^{(i)} = (\tilde{x}_t^{(i)}, \tilde{\theta}_t^{(i)})\}_{i=1}^N$  denote particles jointly generated from  $p(x_t, \theta|y_{1:t-1})$  and  $p(x_t, \theta|y_{1:t})$  respectively. Then the particle approximation of the Bayesian

learning process (Kitagawa (1998)) is given by

$$p(z_t|y_{1:t-1}) \simeq \frac{1}{N} \sum_{i=1}^N p(z_t|\tilde{z}_{t-1}^{(i)}), \quad (8)$$

$$p(z_t|y_{1:t}) \simeq \sum_{i=1}^N w_t^{(i)} \delta(z_t - z_t^{(i)}), \quad w_t^{(i)} = \frac{p(y_t|z_t^{(i)})}{\sum_{i=1}^N p(y_t|z_t^{(i)})}. \quad (9)$$

This is a rather straightforward generalization of the particle filtering.

#### ALGORITHM: PARTICLE LEARNING

**Step 0:** Set the starting values of  $N$  particles  $\{z_0^{(i)}\}_{i=1}^N$ .

**Step 1:** Propagate  $z_t^{(i)}$  from  $p(z_t|\tilde{z}_{t-1}^{(i)})$ ,  $(i = 1, \dots, N)$ .

**Step 2:** Compute weight  $w_t^{(i)} \propto p(y_t|z_t^{(i)})$ .

**Step 3:** Resample  $\{\tilde{z}_t^{(i)}\}_{i=1}^N$  from  $\{z_t^{(i)}\}_{i=1}^N$  with weight  $w_t^{(i)}$ .

Once we generate  $\{\tilde{\theta}_t^{(i)}\}$  by the particle learning, we can treat them as a Monte Carlo sample of  $\theta$  drawn from the posterior density  $p(\theta|y_{1:t})$ . Thus we calculate the posterior statistics on  $\theta$  with  $\{\tilde{\theta}_t^{(i)}\}$  in the same manner as the traditional Monte Carlo method or the state-of-the-art Markov chain Monte Carlo (MCMC) method.

The computational burden of particle filtering will be prohibitively taxing as the number of particles  $N$  increases. The number of the likelihood  $p(y_t|x_t^{(i)})$  to be evaluated, the number of executions for constructing the CDF of particles, and the number of particles to be generated in propagation will increase in  $\mathcal{O}(N)$ . The number of executions for resampling with the CDF will increase in  $\mathcal{O}(N^2)$  when we use a naive resampling algorithm, but it can be reduced to  $\mathcal{O}(N \log N)$  with more efficient algorithms, which we will discuss in the next section. Thus sequential particle-by-particle execution of each step in the particle filtering (and learning) algorithm is inefficient when  $N$  is large, even though the particle filtering method by construction requires a large number of particles to guarantee precision in the estimation.

To reduce the time for computation, we propose to parallelize all steps in particle filtering so that we can execute the parallelized particle filtering algorithm completely inside the GPU. The key to constructing an efficient parallel algorithm is asynchronous out-of-order execution of jobs assigned to each processor. We need to keep a massive number of processors in the GPU

as busy as possible to fully exploit a potential computational power of the GPU. Therefore each processor should waste no milliseconds by waiting for other processors until they complete their jobs. If the order of execution is independent of the end results, asynchronous out-of-order execution is readily implemented. In this situation, parallelization is rather straightforward. In the particle filtering method, this is the case for computation of the likelihood and the propagation step and these steps can be computed in parallel without any modifications. For constructing the CDF and resampling particles, on the other hand, the conventional algorithm does not allow asynchronous out-of-order execution and thus parallelization of these steps can be tricky. In order to devise a fully parallelized particle filtering method, we need to develop new algorithms for parallelization for these steps. In the next section, we describe how to implement the CDF construction and resampling in a parallel computing environment.

### 3 Full Parallelization of Particle Filtering

#### 3.1 Fully parallelized CDF construction

The goal of resampling is to generate  $N$  random integers, which are the indices of particles, from a discrete distribution on  $\{1, \dots, N\}$  with the cumulative distribution function,

$$q(i) \propto \sum_{j=1}^i p(y_t | z_t^{(j)}), \quad (i = 1, \dots, N). \quad (10)$$

Therefore we need to construct the CDF (10) before we perform resampling.

Hendeby et al. (2007) developed an algorithm designed for parallel execution of the CDF construction. The process consists of two procedures: the forward adder and the backward adder. These are illustrated in Table 1. Suppose that we have four particles and their weights are given by  $\{2, 4, 3, 1\}$ . What we want to compute is the cumulative sum,  $\{2, 6, 9, 10\}$ . First we apply the forward adder as described in Panel (a) of Table 1. In the initial step of the forward adder, each weight is assigned to a node (nodes are corresponding to threads in the GPU). Let  $w_j^{(i)}$  denote the sum of weights of Node  $i$  in Step  $j$ . Thus the initial states of the nodes are  $w_1^{(1)} = 2$ ,  $w_2^{(1)} = 4$ ,  $w_3^{(1)} = 3$ , and  $w_4^{(1)} = 1$ . Then in each step of the forward adder, a neighboring pair of nodes is combined to form a new node which inherits the sum of weights from the two combining nodes as its new weight. For example, in Step 2 of Table 1(a), a pair



of nodes,  $w_1^{(1)} = 2$  and  $w_1^{(2)} = 4$ , is combined into a new node  $w_2^{(1)} = 2 + 4 = 6$ . We repeat this procedure until all nodes are collapsed into a single node whose weight is the sum of all weights.

Then we apply the backward adder as described in Panel (b) of Table 1. Note that the numbering of steps is reversed in Table 1 (b) as this is intentional. In the initial step of the backward adder, we start at the last node in the forward adder. Let  $s_j^{(i)}$  denote the sum of weights up to Particle  $i$  in Step  $j$ . In Panel (b) of Table 1, we already have  $s_3^{(4)} = 2 + 4 + 3 + 1 = 10$  by the forward adder. As it moves backward, the current node branches into two nodes in each step. The new right node inherits the same value of the current node, while the partial sum of the new left node equals the value of the current node minus the weight of the right node in the corresponding step of the forward adder. For example, in Step 1 of Table 1(b), the node with  $s_2^{(2)} = 6$  branches to  $s_1^{(1)} = s_2^{(2)} - w_1^{(2)} = 6 - 4 = 2$  and  $s_1^{(2)} = s_2^{(2)} = 6$ . Note that the number of particles should be an integer power of 2 in order to apply this algorithm.

Table 1: Parallel computation of the cumulative sum

(a) Forward adder

Step 1	$w_1^{(1)} = 2$	$w_1^{(2)} = 4$	$w_1^{(3)} = 3$	$w_1^{(4)} = 1$
Step 2	$w_2^{(1)} = 2 + 4 = 6$		$w_2^{(2)} = 3 + 1 = 4$	
Step 3	$w_3^{(1)} = 6 + 4 = 10$			

(b) Backward adder

Step 3	$s_3^{(4)} = 10$			
Step 2	$s_2^{(2)} = 10 - 4 = 6$		$s_2^{(4)} = 10$	
Step 1	$s_1^{(1)} = 6 - 4 = 2$	$s_1^{(2)} = 6$	$s_1^{(3)} = 10 - 1 = 9$	$s_1^{(4)} = 10$

### 3.2 A Review of the Conventional Resampling Algorithms

Before we proceed to describe our parallel resampling algorithm, we briefly review several conventional sequential resampling algorithms. In the most naive manner, the resampling algorithm is expressed as

```
for(i in 1:N) {
```

```

u <- runif(1);
for(j in 1:N) {
  if(u < q[j])
    break;
}
sampled[i] <- j;
}

```

Although this is the most basic and exact resampling procedure, it is extremely time-consuming  $\mathcal{O}(N^2)$  operation, as we need to search out each particle one by one in the whole CDF. One simple way to improve it is by sorting the uniform variates before the search process starts:

```

u <- sort(runif(N));
for(i in 1:N) {
  for(j in i:N) {
    if(u[j] < q[j])
      break;
  }
  sampled[i] <- j;
}

```

Compared to the naive algorithm, this algorithm starts off by sampling all necessary random numbers from the uniform distribution, sorts them in ascending order, and then sequentially resamples particle with them. By sorting the uniform variates, each search for a particle can be started where the last research was left off. The offset of this algorithm is that, although the resampling procedure is a more efficient  $\mathcal{O}(N \log N)$  operation, sorting uniform variates can be computationally strenuous as the number of particles increase, depending on the sorting algorithm we use.

Alternative resampling algorithms have been introduced in order to circumvent the computational strains found in exact resampling algorithms like those above. The stratified resampling algorithm,

```

for(i in 1:N) {

```

```

    u <- (i-1+runif(1))/N; }
    for(j in {i}:N) {
        if(u < q[j])
            break;
    }
    sampled[i] <- j;
}

```

conducts the resampling procedure by generating uniform variates on  $N$  equally spaced intervals  $[(i-1)/N, i/N]$  ( $i = 1, \dots, N$ ). Since only one particle will always be picked for each interval, it does not exactly generate random integers from the the CDF  $\{q(1), \dots, q(N)\}$ . The systematic resampling,

```

u0 <- runif(1);
for(i in 1:N) {
    u <- (i-1+u0)/N;
    for(j in i:N) {
        if(u < q[j])
            break;
    }
    sampled[i] <- j;
}

```

is similar to the stratified resampling but it always chooses an identical point in  $[(i-1)/N, i/N]$  for all  $i$ . In essence, the two alternative resampling algorithms are similar to the aforementioned resampling with sorted uniform variates, but without the time-consuming sorting procedure. However, neither stratified sampling nor systematic sampling is exact. Therefore, they may produce less accurate results compared to exact resampling algorithms. In particular when estimating unknown parameters in particle learning, they may be more problematic because they cannot jointly resample the state variables and the parameters. Thus, despite their computational superiority in a sequential framework, we need to be careful of using them in practice. Additionally, as these methods benefit only from the sequential framework, it is not obvious how to parallelize them.

### 3.3 Fully parallelized resampling

To parallelize the resampling procedure while maintaining its exactness, we have developed a parallel resampling algorithm<sup>2</sup> based on the cut-point method by Chen and Asau (1974).

A *cut-point*,  $I_j$ , for given  $j = 1, \dots, N$  is the smallest index  $i$  such that the corresponding probability  $q(I_j)$  should be greater than  $(j - 1)/N$ . In other words,

$$q(I_j) = \min_{1 \leq i \leq N} q(i) \quad \text{subject to} \quad q(i) > \frac{j-1}{N}, \quad (j = 1, \dots, N). \quad (11)$$

Given cut-points  $\{I_1, \dots, I_N\}$ , random integers between 1 and  $N$  is generated from the CDF  $\{q(1), \dots, q(N)\}$  by the following procedure:

ALGORITHM: CUT-POINT METHOD

**Step 0:** Let  $j = 1$

**Step 1:** Generate  $u$  from the uniform distribution on the interval  $(0, 1)$ .

**Step 2:** Let  $k = I_{\lceil Nu \rceil}$  where  $\lceil x \rceil$  stands for the smallest integer greater than or equal to  $x$ .

**Step 3:** If  $u > q(k)$ , let  $k \leftarrow k + 1$  and repeat **Step 3**; otherwise, go to **Step 4**.

**Step 4:** Store  $k$  as the index of the particle.

**Step 5:** If  $j < N$ , let  $j \leftarrow j + 1$  and go back to **Step 1**; otherwise, exit the loop.

Once all cut-points  $\{I_1, \dots, I_N\}$  are given, parallel execution of the cut-point method is straightforward because the execution of **Step 1** – **Step 3** does not depend on the index  $j$ . The fully parallel resampling algorithm distributed on  $N$  threads is given as follows.

ALGORITHM: PARALLELIZED CUT-POINT METHOD

**Step 0:** Initiate the  $j$ -th thread.

---

<sup>2</sup>Hendeby et al. (2007) developed a parallel resampling algorithm for particle filtering which is specifically designed for GPUs. Their method, however, is dependent on a device specific functionality (rasterization) and its efficiency and scalability is limited by the GPU architecture. Our parallel algorithm, on the other hand, is more versatile and scalable because it requires only basic thread coordination mechanisms such as shared memory and thread synchronization which are provided by most parallel computing systems.

**Step 1:** Generate  $u$  from the uniform distribution on the interval  $(0, 1)$ .

**Step 2:** Let  $k = I_{\lceil Nu \rceil}$  where  $\lceil x \rceil$  stands for the smallest integer greater than or equal to  $x$ .

**Step 3:** If  $u > q(k)$ , let  $k \leftarrow k + 1$  and repeat **Step 3**; otherwise, go to **Step 4**.

**Step 4:** Store  $k$  as the index of the particle.

**Step 5:** Wait until all threads complete the job. Otherwise, exit the loop.

However, the conventional algorithm for computation of the cut-points (see Fishman (1996, p.158) for example) is not friendly to parallel execution. Thus, we have developed an efficient algorithm for parallel search of all cut-points. To devise such a search algorithm, let us introduce

$$L_j = \lceil Nq(j) \rceil, \quad (j = 1, \dots, N).$$

and  $L_0 = 0$  for convention. Due to the monotonicity of the cumulative distribution function, we observe

1.  $0 = L_0 < L_1 \leq \dots \leq L_N = N$ .
2. If  $L_{j-1} < L_j$ , a cut-point such that

$$q(I_k) = \min_{1 \leq i \leq N} q(i) \quad \text{subject to} \quad Nq(i) > k - 1, \quad (k = L_{j-1} + 1, \dots, L_j)$$

is given as  $I_k = j$ .

3. If  $L_{j-1} = L_j$ ,  $j$  is not corresponding to any cut-points.
4.  $L_1 = I_1$  always holds.

The above properties give us a convenient criterion to check whether a particular  $L_j$  is a cut-point or not and it leads to the following multi-thread parallel algorithm to find all cut-points.

ALGORITHM: PARALLELIZED CUT-POINT SEARCH

**Step 0:** Initiate the  $j$ -th thread.

**Step 1:** Compute  $L_j = \lceil Nq(j) \rceil$ .

**Step 2:** Let  $k = L_j$ .

**Step 3:** If  $k > L_{j-1}$ , let  $I_k = j$  and end the thread.

**Step 4:** Let  $k \leftarrow k - 1$  and go to **Step 3**.

With a fully parallel resampling algorithm, particle filtering can be executed in a fully parallel manner without any compromise. Additionally, as the particle filtering algorithm (and the particle learning algorithm) is conducted completely on the GPU and each particle goes through the algorithm on each designated core without syncing, the advantage of parallel computing is gained to the fullest while its shortcoming is kept to its minimum (data transfer between the GPU's device memory and the CPU's host memory only occurs at the beginning and the end of the computation).

## 4 Numerical Experiment

In our experiment, we use a stochastic trend with noise model:

$$\begin{cases} y_t = x_t + \nu_t, & \nu_t \sim \mathcal{N}(0, \sigma^2) \\ x_t = x_{t-1} + \epsilon_t, & \epsilon_t \sim \mathcal{N}(0, \tau^2) \end{cases} \quad (12)$$

as the benchmark model for performance comparison. In (12), we set  $x_0 = 0$ ,  $\sigma^2 = 1$ ,  $\tau^2 = 0.1$ , and generate  $\{y_1, \dots, y_{100}\}$ . Then we treat  $\sigma^2$  and  $\tau^2$  as unknown parameters and apply the particle learning algorithm by Carvalho et al. (2010) to (12). The prior distributions are

$$x_0 \sim \mathcal{N}(0, 10), \quad \sigma^2 \sim \mathcal{IG}(5, 4), \quad \tau^2 \sim \mathcal{IG}(5, 0.4)$$

To demonstrate the effectiveness of our new parallel algorithm, we will compare the following types of algorithms:

- Sequential algorithm on the CPU

**CPU(n):** Naive resampling with single precision

**CPU(s):** Resampling with sorted uniform variates with single precision

- Parallel algorithm on the GPU

***GPU(sp)***: Parallel resampling by the cut-point method with single precision

***GPU(dp)***: Parallel resampling by the cut-point method with double precision

The first two are conventional sequential algorithms for resampling. The code for the parallel algorithm is written in CUDA while that for the sequential algorithms is in C. Both are compiled and executed on the same Linux PC with specifications shown in Table 2. Alternative resampling algorithms, such as the stratified and systematic resampling, are not considered here, as they are not exact resampling. However, if we exclude the time consumed by the sorting procedure from the resampling time of *CPU(s)*, we get a very good estimate of how long they might take.

Table 2: Hardware specifications

	CPU (Intel Core-i7 2700K)	GPU (NVIDIA GTX580)
Core clock rate	3.50GHz	772MHz
Number of cores	4	512
Memory	8GB	3GB

For each algorithm, we execute the particle learning ten times with the same generated path,  $\{y_1, \dots, y_{100}\}$ , and recorded the execution time of each trial. To avoid the influence of possible outliers, we took the average of the middle five of them. The results are listed in Table 3 and the plots of the total execution time against the number of particles are shown in Figure 1.

The results clearly show that our new parallel algorithm, which runs completely in parallel and keeps all executions within the GPU, can be extremely effective compared to conventional sequential algorithms. As the number of particles increases (and the precision of the estimate increases), *GPU(sp)* outperforms that of *CPU(n)* by more than 200x in the case of 1,048,576(=  $2^{20}$ ) particles and 1,671x in the case of 8,388,608(=  $2^{23}$ ) particles. Even when compared with the exact resampling with sorted uniform variates (*CPU(s)*), *GPU(sp)* is consistently faster by more than 20x when the number of particles is 131,072 or more. In the comparison between *GPU(sp)* and *GPU(dp)*, the difference is somewhere around two times, which is consistent with intuition. Interestingly, the computation on the GPU in double precision is still a good 5-10x faster than that of the CPU in single precision, which demonstrates the sheer power of

Table 3: Comparison in execution time

	Number of particles					
Time	1,024	4,096	16,384	131,072	1,048,576	8,388,608
(i) $CPU(n)$	30	80	390	7,050	307,600	19,855,870
(ii) $CPU(s)$	48	188	754	6,088	49,266	330,742
(iii) $GPU(sp)$	17	21	44	215	1,526	11,881
(iv) $GPU(dp)$	20	25	129	436	—	—
Ratio	1,024	4,096	16,384	131,072	1,048,576	8,388,608
(i) / (iii)	1.7	3.8	7.7	32.7	201.5	1671.3
(ii) / (iii)	2.8	9.0	17.3	28.3	32.3	27.8
(ii) / (iv)	2.4	7.4	5.8	14.0	—	—
(iv) / (iii)	1.1	1.2	3.0	2.0	—	—

Note: the values of execution time are in milliseconds.




Figure 1: Plots of execution time against the number of particles



parallel processing on the GPU. Due to memory failure,  $GPU(dp)$  failed for particles more than a million, though this could be remedied by upgrading the GPU to the one with more memory or using multiple GPUs.

To see which part of the particle learning contributes to the reduction in execution time, we divide the cycle of particle learning into the following steps;

***Initialize:*** set the starting values of the particles;

***CDF:*** compute the likelihood and construct the CDF;

***Resample:*** resample the particles with the CDF;

***Propagate:*** propagate a new set of particles;

***Store:*** store the generated particles into the CPU's host memory (GPU only);

***Other:*** keep the results and proceed with the particle learning;

The results in the case of 131,972 particles are listed in Table 4. The tendency we observe in Table 4 is similar in the other cases.

Breaking down the execution time gives us deep insights into how the GPU architecture works and its strong and weak points. Examining the results in  $CPU(s)$ , we first notice that the *CDF* step and *Propagation* step put together occupy the bulk of the total execution time, while the *Resampling* step only accounts for less than ten percent of the total execution time and much of it coming from the sorting step. Looking closely into the gains by parallelization in each step, the largest comes from the *CDF* step with a gain of 248x, followed by the *Propagation* step with a gain of 45.3x, then followed by the *Resampling* step with a gain of 11.9x. Although the gain in the *Resampling* step has less of an impact compared with the overwhelming gain in *CDF* and *Propagation*, it does not overshadow the fact that it gained 2.7x in single precision even if we ignore the time spent in sorting the uniform random variates and focus on the resampling procedure only. That implies that our parallel resampling on the GPU can beat the stratified resampling on the CPU since the stratified resampling is roughly equivalent to the resampling with sorted uniform variates without sorting in terms of computational complexity. As for the *Other* step,  $CPU(s)$  and  $GPU(sp)$  is identical. This is because for both algorithms, all executions of this step are conducted only on the CPU. Thus, we observe no difference. Finally,

Table 4: Breakdown of execution time

		Cycle of particle learning					
Time	Initialize	CDF	Resample	Propagate	Store	Other	Total
(i) <i>CPU(s)</i>	26	2652	454 (102)	2900	—	56	6088
(ii) <i>GPU(sp)</i>	0.72	11	38	64	46	56	215
(iii) <i>GPU(dp)</i>	1.51	19	61	174	89	92	436
Ratio	Initialize	CDF	Resample	Propagate	Store	Other	Total
(i) / (ii)	35.9	248.0	11.9 (2.7)	45.3	—	1.0	28.3
(iii) / (ii)	2.1	1.8	1.6	2.7	1.9	1.7	2.0

Notes: (a) the number of particle is 131,072;

(b) the values of execution time are in milliseconds;

(c) the number in parentheses is corresponding to the time excluding the sorting step.

we observe a good amount of reduction in initiating the particle learning algorithm by our parallel algorithm; however, the time spent in initiation is quite trivial, in particular when the number of the sample period  $T$  is large ( $T = 100$  in our experiment).

Although it is clear that our parallel algorithm is superior to the conventional sequential algorithm through every step, Table 4 indicates that there is one drawback of using the GPU. That is memory transfer. The *Store* step measures the time it takes to transfer the generated particles from the GPU’s device memory to the CPU’s host memory. Table 4 shows that it takes up roughly 15-20% of the execution time. Note that, for fairness of the experiment, the GPU returns all of the particles it generated to the CPU’s host memory. If we were to return only the mean, the variance, and other statistics of the state variables and parameters, the time for the *Store* step can be cut down significantly.

## 5 Conclusion

In this study, we have developed a new algorithm to perform particle filtering and learning in a parallel computing environment, in particular on GPGPUs. Our new algorithm has several advantages. First, it enables us to keep all executions of the particle filtering (and learning) algorithm within the GPU so that data transfer between the GPU’s device memory and the CPU’s host memory is minimized. Second, unlike the stratified sampling or the systematic sampling, our parallel sampling algorithm based on the cut-point method can resample particles exactly from their CDF. Lastly, since our algorithm does not utilize any device specific functionalities, it is straightforward to apply our algorithm to a multiple GPU system or a large grid computing system.

Then we conducted a Monte Carlo experiment in order to compare our parallel algorithm with conventional sequential algorithms. In the experiment, our algorithm implemented on the GPU yields results far better than the conventional sequential algorithms on the CPU. Although we keep the SSM as simple as possible in the experiment, our parallel algorithm can also be applied to more complex models without any fundamental modifications to the programming code and this little investment will return a significant gain in execution time instantaneously.

Our fully parallelized particle filtering algorithm is beneficial for various applications that require estimating powerful but complex models in a shorter span of time; ranging from motion

tracking technology to high-frequency trading. We even envision that one can perform real-time filtering of the state variables and the unknown parameters in a high-dimensional nonlinear non-Gaussian SSM on an affordable parallel computing system in a completely parallel manner. That would pave the way for a new era of computationally intensive data analysis.

## Reference

- [1] Carvalho, C. M., Johannes, A. M., Lopes, H. F., and Polson, N. G. (2010), “Particle Learning and Smoothing,” *Statistical Science*, 25, 88–106.
- [2] Chai, X. and Yang, Q. (2007), “Reducing the calibration effort for probabilistic indoor location estimation,” *IEEE Transactions on Mobile Computing*, 6, 649–662.
- [3] Chen, H. C., and Asau, Y. (1974), “On generating random variates from an empirical distribution,” *American Institute of Industrial Engineers (AIIE) Transactions*, 6, 163–166.
- [4] Douc, R., Cappe, O., and Moulines, E. (2005), “Comparison of Resampling Schemes for Particle Filtering,” *Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis*.
- [5] Doucet, A., and Johannes, A. M. (2011), “A tutorial on particle filtering and smoothing: fifteen years later,” *The Oxford Handbook of Nonlinear Filtering*, Oxford University Press.
- [6] Dukic, V. M., Lopes, H. F., and Polson, N. G. (2009), “Tracking flu epidemics using Google flu trends and particle learning,” *Working paper*, Univ. Chicago Booth School of Business.
- [7] Durham, G. B., and Geweke, G. (2011), “Massively Parallel Sequential Monte Carlo for Bayesian Inference,” *Working paper*.
- [8] Fearnhead, P. (2002), “Markov chain Monte Carlo, sufficient statistics, and particle filters,” *J. Comput. Graph. Statist.* 11 848–862.
- [9] Fishman, G. S. (1996), “Monte Carlo: Concepts, Algorithms, and Applications,” Springer.
- [10] Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993), “Novel approach to nonlinear/non-Gaussian Bayesian state estimation,” *IEEE Proceedings-F*, 140, 107–113.
- [11] Harvey, A. C., Ruiz, E., and Shephard, N. (1994), “Multivariate stochastic variance models,” *Review of Economic Studies*, 61, 247–264.
- [12] Hendeby, G., Karlsson, R., Gustafsson, F. (2010), “Particle filtering: the need for speed,” *Journal on Advances in Signal Processing*, 22.

- [13] Hendeby, G., Hol, J. D., Karlsson, R., and Gustafsson, F. (2007), “A graphics processing unit implementation of the particle filter,” *Proceedings of the 15th European Statistical Signal Processing Conference (EUSIPCO '07)*, 1639–1643, Poznan, Poland.
- [14] Johannes, M. and Polson, N. G. (2008), ”Exact particle filtering and learning”. *Working paper*, Univ. Chicago Booth School of Business.
- [15] Johannes, M., Polson, N. G. and Yae, S. M. (2008), “Non-linear filtering and learning,” *Working paper*, Univ. Chicago Booth School of Business.
- [16] Kalman, R. E. (1960), “A new approach to linear filtering and prediction problems,” *Transactions of the ASME, Journal of Basic Engineering*, 82, 35–45.
- [17] Kitagawa, G. (1996), “Monte Carlo filter and smoother for non-Gaussian nonlinear state space models,” *Journal of Computational and Graphical Statistics*, 5, 1–25.
- [18] Kitagawa, G. (1998), “A self-organizing state-space model,” *Journal of the American Statistical Association*, 93, 1203–1215.
- [19] Lee, A., Yau, C., Giles, M. B., Doucet, A., and Holmes, C. C. (2010), “On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods,” *Journal of Computational and Graphical Statistics*, 19, 769–789.
- [20] Liu, J. and West, M. (2001), “Combined parameters and state estimation in simulation-based filtering,” *Sequential Monte Carlo Methods in Practice* (A. Doucet, N. de Freitas and N. Gordon, eds.). New York: Springer-Verlag, 197–223.
- [21] Lopes, H. F. and Tsay, R. S. (2011), “Particle filters and Bayesian inference in financial econometrics,” *Journal of Forecasting*, 30, 168–209.
- [22] Maskell, S., Alun-Jones, B., and Macleod, M. (2006), “A single instruction multiple data particle filter,” *Proceedings of Nonlinear Statistical Signal Processing Workshop (NSSPW '06)*, Cambridge, UK.
- [23] Mihaylova, L., Angelova, D., Honary, S., Bull, D. R., Canagarajah. C. N., and Ristic, B. (2007), “Mobility tracking in cellular networks using particle filtering,” *IEEE Transactions on Wireless Communications*, 6, 3589–3599.

- [24] Montemayor, A. S., Pantrigo, J. J., Sanchez, A., and Fernandez, F. (2004), “Particle filter on GPUs for real time tracking,” *Proceedings of the International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '04)*, 94, Los Angeles, CA, USA.
- [25] Montemerlo, M., Thrun, S., Koller, D., and Wegbreit, B. (2003), “FastSLAM 2.0: an improved particle filtering algorithm for simultaneous localization and mapping that provably converges,” *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, 1151–1157, Acapulco, Mexico. *Computational Statistics & Data Analysis*, 56, 3690–3704.
- [26] Pitt, M. and Shephard, N. (1999), “Filtering via simulation: Auxiliary particle filters,” *J. Amer. Statist. Assoc.*, 94, 590–599.
- [27] Polson, N. G., Stroud, J. and Muller, P. (2008), “Practical filtering with sequential parameter learning,” *J. Roy. Statist. Soc. Ser. B*, 70, 413–428.
- [28] Storvik, G. (2002), “Particle filters in state space models with the presence of unknown static parameters,” *IEEE Trans. Signal Process*, 50, 281–289.
- [29] West, M. and Harrison, J. (1997), *Bayesian Forecasting and Dynamic Models*, 2nd ed. Springer.
- [30] Zou, Y. and Chakrabarty, K. (2007), “Distributed mobility management for target tracking in mobile sensor networks,” *IEEE Transactions on Mobile Computing*, 6, 872–887.